

Building an Information Management Infrastructure in the 90s: The Vanderbilt Experiment

William W. Stead, M.D., Ruby Borden, R.N., Paul McNulty, Dean F. Sittig, Ph.D.

Center for Biomedical Informatics, Vanderbilt University,
1500 21st Avenue South, Suite 2000, Nashville, Tn. 37232-8143

ABSTRACT

The course that an organization takes to create a competitive information management infrastructure is determined by a series of decisions, each of which balances tradeoffs. Key success factors include sequencing projects to reflect data requirements; obtaining benefits as cost is incurred; establishing an architecture that permits integration of applications; managing project scope; and establishing a data friendly culture.

INTRODUCTION

What course should an organization take once it has identified the need for a competitive information management infrastructure? This question is being asked with increasing frequency by organizations of all types and sizes as they look for ways to compete in the changing health care marketplace. Different alternatives cannot be labeled clearly as right or wrong. The correct course is determined by a series of decisions, each of which balances trade-offs. For example, the speed with which a pressing need must be met determines the degree to which the solution can be state-of-the-art or a platform for future development.

This paper discusses the choices that were made during the first two years of the effort to establish the new information management infrastructure at Vanderbilt University Medical Center (VUMC). Emphasis is placed upon the thought processes and strategies which led to those choices. These concepts should be useful to other institutions even though they may make different selections to reflect the changing technologic environment, their institutional culture, management priorities, or personal biases.

BACKGROUND

VUMC is an integral part of Vanderbilt University and consists of the School of Medicine, the School of Nursing, and the Vanderbilt University Hospital and

Clinics. The School of Medicine has 822 full-time faculty members in 28 departments. The School of Nursing has a faculty of 49 and awards graduate degrees. The Vanderbilt University Hospital is currently licensed as a 661-bed facility. During FY 92 there were 27,722 discharges and 345,000 visits to the outpatient facilities.

In July, 1991, the leadership of VUMC made the strategic decision to create a competitive edge for the institution in the area of management and utilization of information as a strategic resource [1]. That decision led to a number of related initiatives. First, a Center for Biomedical Informatics was established within the School of Medicine to allow linkage of the research and training components of academic informatics with provision of the operation and decision support systems that underpin the institution. Second, VUMC began construction of a new biomedical library designed to test models of the library of the future. Third, we began implementation of a new information management infrastructure, starting with an enterprise backbone network and a patient care information system.

At the start of the effort, the core VUMC information systems included: HBO's Medipac admission/discharge/transfer (ADT), scheduling, chart tracking, and billing packages operating on an IBM ES/9000; the DECrad radiology application operating on a VAX 11/785; CHC's laboratory application running on a Stratus XA2000. Rudimentary interfaces were in place between these systems to exchange ADT transactions and billing information. A number of independent Novell LANs were in place. Several were daisy-chained together, and some were connected to a campus broadband Ethernet backbone.

SEQUENCING PROJECTS TO REFLECT DATA REQUIREMENTS

At the beginning of the effort, we intended to start by building a robust backbone network. Then, the

existing systems would be attached to this network, and network attached workstations would be utilized as a single point of access to their data. This approach was modeled after the course taken at Columbia Presbyterian Medical Center [2]. It was intended to buy the time necessary to re-engineer applications so as to provide an integrated patient database as a background task.

Early in the effort, we came under intense pressure to update the patient accounting software, particularly in the areas of professional billing, electronic data exchange, and work-flow automation. Although the necessary application software could have been purchased, the requirements analysis pointed out that the new applications would not be effective. The data available from the institution's front-end systems were inadequate to support sophisticated back-end applications.

The understanding that developed during this analysis is generalized in Figure 1. The concept is that systems should be put in place in a sequence that allows the first system to establish a foundation of data for the next system. Examination of data dependencies suggests that systems should be installed in three layers. The first layer automates various processes that underpin the institution. These systems form a foundation because they are the logical place to capture data about patients. They must be implemented in a manner that guarantees the accuracy and non-ambiguity of that data. The second layer of system involves databanking, either in the form of computer-based records for a patient or outcome/utilization databases. These systems integrate data from a variety of sources. Systems that databank cannot be put in place until the necessary feeder systems are up and running. The third layer of system involves intelligent decision support and process control. Implementation of such systems should begin after the first two layers are in place because the databanks are resources from which knowledge should be derived to construct the intelligent logic modules, and the data passing through the process automation systems is the best trigger for activation of the intelligent modules.

Even within each of these layers, the systems should be implemented in a sequence based upon data dependencies. For example, ADT and clinic processing systems are a first step because they establish patient identification and episodes of care. Orders should precede automation of ancillary tasks and result reporting to permit result updates to be

handled in a non-ambiguous fashion. Certain systems may be implemented in parallel. For example, once order capture is in place, billing can be brought up concurrently with result reporting.

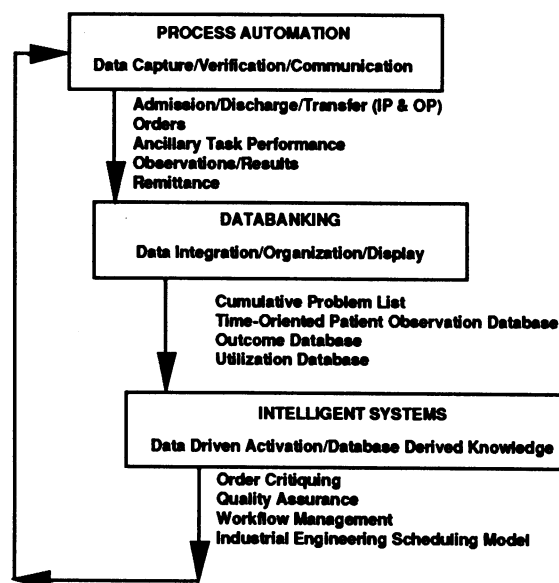


Figure 1. Foundation of Data Concept: Sequencing system projects so that each layer provides a data foundation for the next.

The foundation of data concept is important for two reasons. First, the users press for the type of functionality provided by the second or third layers of systems. An organization's chief information architect faces a difficult challenge in convincing people to be patient while the first layer of systems is installed or modified to provide quality data. If the organization implements systems out of sequence, data entry requirements increase and data accuracy decreases. Second, clear prior forethought is required during the implementation of the first layer of process automation systems if they are to also function as feeder systems for the subsequent layers.

After the foundation of data concept was explained, VUMC agreed to delay installation of new patient accounting software. Instead, we decided to implement new front-end systems in parallel with implementation of the backbone network.

OBTAINING BENEFIT AS COST IS INCURRED

The aggregate cost of a state-of-the-art, full function, information management infrastructure is beyond the realistic reach of any organization if it is treated as a single expense. For example, rewiring costs alone for an institution such as Vanderbilt

would exceed \$20 million. It is difficult to justify tying up that amount of capital without an immediate commensurate benefit. VUMC's decisions about how to build its backbone data network point out several strategies for controlling up-front cost while building for the future.

We began by putting in place an intrabuilding wiring standard and mandating its use in all new construction or those renovations which involved construction of new walls. Over time, this standard will allow us to take advantage of the ongoing growth and evolution of the Medical Center to rewire the majority of the physical plant without a dedicated project.

The intrabuilding wiring standard, and the plans for the interbuilding network and the intrabuilding riser systems, were developed so as to treat the communication spaces, the transmission media, and the network electronics as separate components. This approach reflected the fact that the best approach to networking today will be out-of-date tomorrow. It is more important for the design to permit evolution than to be perfect at any point in time. We attempted to make the communication spaces robust enough to last the life of the building, whereas we guessed that the transmission media would need to be updated in five to ten years and the network electronics within three to five years. Our decision to design all of the cable trays, conduits, and outlet boxes to accommodate fiber optic cable, while installing only copper cable in many places to save money, reflected this approach. Similarly, where we did install fiber optic cable, we installed only the multi-mode fiber that is needed today rather than both single and multi-mode fiber. At the same time, we placed a spare interduct so that it would be easy to come back and install the single mode fiber when the technology becomes available to use it fully.

We had to decide when to allow the implementation to follow behind the identified need, and when we should get out in front of that need. The former course limits idle capital investment, but the latter course is necessary to stimulate creativity and to increase the rate at which new need is identified. VUMC's existing network traffic could easily have been handled by either the Ethernet or Token Ring protocols. However, we wanted to stimulate experimentation with image-based applications which would require a higher bandwidth. We elected to install a FDDI ring to connect the five Medical Center buildings that had the highest network

traffic. The intrabuilding riser system is based upon the less expensive Token Ring technology. The eight buildings with lower volumes of network traffic were attached by Token Ring to the riser system of the closest building on the FDDI ring. The cable plant includes spare fibers that can be used to connect a workstation in any building directly to the FDDI ring. This approach allows us to quickly, on a case-by-case basis, deliver a high bandwidth FDDI pathway to any user interested in a pilot project without incurring the cost of a complete FDDI implementation.

IMPLEMENTING APPLICATIONS

An institution that is ready to install a new patient care information system faces a dilemma. Proven, off-the-shelf systems that provide an integrated patient database were architected in the 1970s. Systems built around state-of-the-art technical approaches, either do not provide an integrated patient database, or are not out of the alpha and beta testing stages [3]. We felt that the need for a new front-end patient care system was too great to risk the delays that result from developmental projects involving large applications.

We adopted the information management architecture developed at Duke University as our model for permitting distributed applications while maintaining an integrated patient database [4]. This architecture is designed to disassociate the information system into components; specifically separating the application programs which provide end-user function from the underlying databases. Master databases of items such as patient identification or data item definitions are to be maintained as resources independent of any application package. Distributed applications pass information of general interest to a central transaction database/data hub which is in turn responsible for getting the information to other systems that need to be updated and to the institution's integrated data repository.

We elected to purchase Shared Medical System's (SMS) Invision software to support the front-end patient care information system functions of ADT, clinic processing, order entry, and result reporting. In making this selection, we focused upon the technical aspects of the software. The way in which it balances the tradeoff between general purpose tools and efficient utilization of processing and data storage was a significant deciding factor. SMS's track record of supporting evolution instead of revolution was

equally important. We looked closely at the end-user function provided by those elements of the system which we were going to install in the first two phases. We reduced confusion by not examining capabilities of other elements of the system. We realized that it would be at least two years before we could consider implementing any of those elements and that their capability, together with that of their competitors, would have changed markedly by that time. Given our architectural direction, we will be able to support applications developed by multiple vendors, if necessary, to provide subsequent layers of end-user function.

Despite the complexity of the VUMC environment, we decided to install the Invision product in "vanilla" form. This decision meant that we would not change the source code provided by SMS. In most cases, the architectural tools provided with the system permitted us to meet our needs within the context of the Invision product. Where that was not the case, rather than modifying a piece of Invision, we turned off the offending portion and substituted a complete module developed and maintained by Vanderbilt. If exercised sparingly, this approach allows flexibility in essential areas, while permitting the vendor to update and maintain their product independently from VUMC.

We balanced our decision to minimize risk by installing a stable application package, with a decision to move ahead in terms of technology platforms. Specifically, we decided to use an FDDI backbone network and IBM PS/2 Model 90's (Intel 486 workstations), running the OS/2 2.0 operating system, as the access devices.

We decided to develop four elements of the system at VUMC. First, we built our own OS/2 desk-top and sign-on processor because we wanted to position ourselves to provide seamless access to the complete set of VUMC systems, not just those provided by SMS. Second, we elected to build a generic interface subsystem to sit between Invision and the other VUMC systems. This subsystem is the first step toward the communications hub that is an integral part of our long-term architectural direction. Any data item generated by Invision is passed to that interface subsystem, and it in turn controls what is passed to each of the other systems. This approach lets VUMC change the data requirements of its other systems without interfering with SMS. Third, information about VUMC's environment or processes, which would otherwise be contained in

hard-coded Invision screens, are being maintained in VUMC built and controlled relational database (IBM/DB2) tables. This approach increases the maintainability of the Invision application and provides direct access to that information by other systems. Fourth, we have elected to build our own application to manage physician selection of orders. We are again using the relational tables of orderable items to provide various clinically relevant pathways without introducing an Invision maintenance problem.

MANAGING PROJECT SCOPE

Although the VUMC effort to implement a full function information management infrastructure will extend over a five to seven year period, each part of the effort is being phased so that no single step takes longer than six to eighteen months. This approach achieves several objectives. First, benefit begins to accrue early in the effort, and it is therefore easier to defend the increase in expenditures that will be necessary to achieve our goal. Second, mistakes that are identified in the early stages can be corrected in the design process for subsequent stages. Finally, if circumstances change and the effort must be stopped before completion, the elements that have made it to production will continue to be used, justifying their cost.

For phased implementation to work, each phase must stay roughly on schedule. Four ground rules are helpful. First, the implementation schedule should be preserved, if necessary, by delaying non-critical functionality. Second, purchased software should be installed in "vanilla" form. Third, those functions, which are part of a phase, should be implemented house-wide. Finally, hospital and clinic operating procedures should be standardized and allowed to adapt to the software.

These rules may seem arbitrary, but they are necessary to implement a project in a reasonable period of time. Flexibility comes from the fact that there are to be subsequent phases. Therefore, items which must be dropped from a phase to meet the schedule can be added back later. Similarly, inconvenience that results from a standard operating procedure can be corrected once experience has demonstrated where variability is essential.

ESTABLISHING A DATA FRIENDLY CULTURE

Information systems are installed to meet a

pressing need, such as to speed up the generation of bills or increase the availability of laboratory results. Data processors and users alike have a tradition of manipulating data definitions and data content to get the job done. They do not realize that they frequently make the data meaningless in the process. Several implementation strategies are helpful in creating a data friendly culture. Data descriptors need to be defined and managed outside of any one system and they need to be standardized across systems. Second, redundant data entry should be eliminated. Third, data accuracy should be approached using total quality management principles in which errors are corrected at the source. Finally, data should be entered directly into the system at the time, place and by the person generating it.

In addition, the system development staff should observe three rules. First, each concept should be represented in a separate field in the database. Second, constraints that are imposed upon data representation by pre-existing systems should be handled by mappings internal to the inter-system interfaces or by cross reference tables. Third, data entry applications should be designed so that questions will not be asked unless the user can be expected to know the correct answers.

CURRENT VUMC STATUS AND NEXT STEPS

VUMC is two years into the effort to establish a competitive information management infrastructure. The interbuilding backbone network is operational and it reaches each of VUMC's thirteen buildings. The intrabuilding riser system has been installed in the hospital and clinic buildings. Elsewhere, interested individuals can connect to the backbone by running a copper cable to the building hub at a cost of \$125-\$500. The total expenditure related to the backbone has been \$1.5 million. The rest of the intrabuilding riser system will be built out, on as needed basis, over the next two years at an expected cost of \$2 million.

The clinic processing and ADT portions of Invision were implemented house-wide in an elapsed time of eight months from the day that tapes were received from SMS. One situation was identified that required a change in source code. Physician order entry, development of a longitudinal patient database,

and replacement of the patient accounting system are the next steps. The order entry effort should be complete at the three year marker and the replacement patient accounting system is targeted for the fourth year of the effort. The development of the longitudinal patient database will parallel those efforts.

CONCLUSION

Establishing an information management infrastructure remains a daunting task in the 1990s. Any plan will require the balancing of trade-offs. No plan will get functionality to the users as fast as it is wanted and needed. The challenge for information architects is to demonstrate that they know where an organization needs to go and that they will get it there over time. The principles described in this paper should be helpful in meeting those challenges.

REFERENCES

- [1]. Stead WW, Baker W, Harris TR, Hodges TM, Sittig DF: A Fast Track to IAIMS: The Vanderbilt University Strategy. Proc 16th Symp Computing Applications Med Care, ed. Frisse ME, McGraw-Hill Inc. 1993:527-531.
- [2]. Clayton PD, Sideli RV, Sengupta S: Open Architecture and Integrated Information at Columbia-Presbyterian Medical Center. MD Computing 1992;9:297-303.
- [3]. Dorenfest SL: History and Impediments to Progress in the Development and Implementation of the Computerized Patient Record. Proc Am Hosp Assoc Annual HIMSS Conf 1993;2:83-92.
- [4]. Stead WW, Borden RB, Boyarsky MW, Crow DS, Mears TP, Stone AA, Woods PJ: A System's Architecture Which Dissociates Management of Shared Data and End-User Function. Proc 15th Symp Computing Applications Med Care, ed. Clayton PD, IEEE 1991:475-480.

ACKNOWLEDGEMENT

This research was supported in part by National Library of Medicine Grant GO8 LM05443, awarded by the National Institutes of Health, Department of Health and Human Services.